

05/19/99



JCS53 U.S. PTO

Please type a plus sign (+) inside this box [+]

PTO/SB/05 (12/97)

Approved for use through 09/30/00. OMB 0651-0032

Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

UTILITY PATENT APPLICATION TRANSMITTAL

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Attorney Docket No. 82771.P286Total Pages 3First Named Inventor or Application Identifier Heather AchillesExpress Mail Label No. EL371007036US

JCS11 U.S. PTO

09/314566



05/19/99

ADDRESS TO: Assistant Commissioner for Patents
 Box Patent Application
 Washington, D. C. 20231

APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents.

1. X Fee Transmittal Form
(Submit an original, and a duplicate for fee processing)
2. X Specification (Total Pages 21)
(preferred arrangement set forth below)
 - Descriptive Title of the Invention
 - Cross References to Related Applications
 - Statement Regarding Fed sponsored R & D
 - Reference to Microfiche Appendix
 - Background of the Invention
 - Brief Summary of the Invention
 - Brief Description of the Drawings (if filed)
 - Detailed Description
 - Claims
 - Abstract of the Disclosure
3. X Drawings(s) (35 USC 113) (Total Sheets 5)
4. X Oath or Declaration (Total Pages 4)
 - a. Newly Executed (Original or Copy)
 - b. Copy from a Prior Application (37 CFR 1.63(d))
(for Continuation/Divisional with Box 17 completed) (**Note Box 5 below**)
 - i. DELETIONS OF INVENTOR(S) Signed statement attached deleting inventor(s) named in the prior application, see 37 CFR 1.63(d)(2) and 1.33(b).
5. Incorporation By Reference (useable if Box 4b is checked)
The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.
6. Microfiche Computer Program (Appendix)

12/01/97

- 1 -

PTO/SB/05 (12/97)

7. ☐ Nucleotide and/or Amino Acid Sequence Submission
(if applicable, all necessary)
a. ☐ Computer Readable Copy
b. ☐ Paper Copy (identical to computer copy)
c. ☐ Statement verifying identity of above copies

ACCOMPANYING APPLICATION PARTS

8. ☐ Assignment Papers (cover sheet & documents(s))
9. ☐ a. 37 CFR 3.73(b) Statement (where there is an assignee)
☒ b. Power of Attorney
10. ☐ English Translation Document (if applicable)
11. ☐ a. Information Disclosure Statement (IDS)/PTO-1449
☐ b. Copies of IDS Citations
12. ☐ Preliminary Amendment
13. ☒ Return Receipt Postcard (MPEP 503) (Should be specifically itemized)
14. ☐ a. Small Entity Statement(s)
b. Statement filed in prior application, Status still proper and desired
15. ☐ Certified Copy of Priority Document(s) (if foreign priority is claimed)
16. ☒ Other: Declaration: Unsigned
Attorney Signature Page, Express Mail Certificate, and Copy of Postcard.

17. If a **CONTINUING APPLICATION**, check appropriate box and supply the requisite information:
☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP)
of prior application No: ☐

18. Correspondence Address

☐ Customer Number or Bar Code Label _____
(Insert Customer No. or Attach Bar Code Label here)
or

☒ Correspondence Address Below

NAME Thomas A. Hassing
BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

ADDRESS 12400 Wilshire Boulevard
Seventh Floor

CITY Los Angeles STATE California ZIP CODE 90025-1026

Country U.S.A. TELEPHONE (408) 720-8598 FAX (408) 720-9397

12/01/97

- 2 -

PTO/SB/05 (12/97)

Approved for use through 09/30/00. OMB 0651-0032
Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

UNITED STATES PATENT APPLICATION

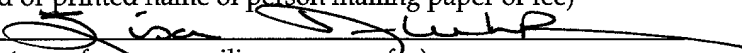
for

METHOD AND APPARATUS FOR DISCARDING DATA PACKETS
THROUGH THE USE OF DESCRIPTORS

Applicants:

Heather Achilles
Terry Pearson
Paul Gallo
Larry Manor

prepared by:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN
12400 Wilshire Boulevard
Los Angeles, CA 90026-1026
(408) 720-8598EXPRESS MAIL CERTIFICATE OF MAILING"Express Mail" mailing label number EL 371 007 036 US
Date of Deposit May 19, 1999I hereby certify that this paper or fee is being deposited with the United States Postal Service
"Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above
and is addressed to the Commissioner of Patents and Trademarks, Washington, D.C. 20231.Lisa F. Luhn
(Typed or printed name of person mailing paper or fee)
(Signature of person mailing paper or fee)

METHOD AND APPARATUS FOR DISCARDING DATA PACKETS THROUGH THE USE OF DESCRIPTORS

FIELD OF THE INVENTION

This invention relates generally to computer networks, and more
5 particularly, to a method and apparatus for discarding data packets through
the use of descriptors.

BACKGROUND OF THE INVENTION

In the field of data routing in computer networks, an Internet service
provider (ISP) user typically has much more stringent requirements than an
10 enterprise user because the routers will be subjected to the adverse Internet
routing environment in the world. There are three typical architectural
requirements that such routers must support, described below.

A. Stable Operation. Although it sounds trivial, the notion of stable
operation has been elusive in the ISP community, as witnessed by various
15 Internet "brown-outs" since it's inception. One paper on Internet scaling
"Scaling the Internet during the T3 NSFNET Years", C. Villamizar, October
22, 1997, articulates the basic requirements which ISPs demand from their
networking equipment in order to provide a stable network. In addition to
forwarding performance and scaling requirements, ISPs typically expect
20 several operational attributes, given below.

1. Stability under adverse conditions. The router must remain
stable and deterministic under arbitrarily high traffic loads or a flood of
routing update changes.

2. Low packet loss to stable destinations. The effects of unstable
25 routes (flapping) should not impact a router's ability to forward traffic to
stable routes.

3. Reasonable fairness and congestion control. Sufficient buffering capacity, avoidance of head-of-line blocking, advanced queuing algorithms, and sophisticated discard techniques must be provided.

5 B. Service Differentiation. Recently it has become clear that service providers cannot make adequate margins by offering flat-rate access and undifferentiated service. The ability to offer tiered services, and to guarantee service levels, is crucial to the economic and competitive health of ISPs. The airline industry's first-class, business-class and coach-class offerings provide a meaningful analogy for Internet service differentiation: a small number of
10 customers are willing to pay for premium service, if it can be guaranteed. The concentrator's must enable ISPs to offer differentiated services based on multiple queues and advanced, intelligent Traffic Management features.

C. Superior Reliability. ISP routers must provide a greater level of reliability and availability than known router architectures. Part of this flows
15 from designing with stability in mind, but providing additional fault tolerance features adds another dimension of resiliency. ISP routers should be designed without any single points of failure, and all software designs should incorporate fault isolation principles.

Therefore, there is a need for a way to route data in computer networks
20 that provides stable operation, service differentiation, and superior reliability. Such an invention should be stable under adverse conditions, insure low packet loss to stable destinations, and provide reasonable fairness and congestion control.

SUMMARY OF THE INVENTION

The present invention provides a method, apparatus and article of manufacture for discarding a data packet. The data packet is classified according to a type of service (TOS) indicator and modified with an internal service class (ISC) indicator according to the TOS indicator. The data packet is modified with a watermark (WM) indicator according to the availability of a system resource. The ISC, WM and a drop preference (DP) indicator of the data packet are compared to a committed information rate (CIR). The packet is discarded if the DP exceeds the CIR.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example and may be better understood by referring to the following description in conjunction with the accompanying drawings, in which like references indicate similar elements and in which:

FIG. 1 shows a block diagram of an L3 card forwarding path compatible with the present invention;

FIG. 2 shows a buffer pool division compatible with the present invention;

FIG. 3 shows an inbound descriptor format compatible with the present invention;

FIG. 4 shows a flow diagram of a method for discarding and forwarding data packets compatible with the present invention;

FIG. 5 shows a flow diagram of a method for discarding data packets compatible with the present invention.

DETAILED DESCRIPTION OF AN EMBODIMENT OF THE INVENTION

In the following description of an embodiment, reference is made to the accompanying drawings which form a part hereof, and in which is shown by way of illustration a specific embodiment in which the disclosed
5 technology may be practiced. It is to be understood that other embodiments may be utilized and structural changes may be made without departing from the scope of the disclosed technology. An embodiment of the disclosed technology, described below, enables a remote computer system user to execute a software application on a network file server.

10 The disclosed technology provides a method, apparatus and article of manufacture for discarding a data packet. The data packet is classified according to a type of service (TOS) indicator and modified with an internal service class (ISC) indicator according to the TOS indicator. The data packet is modified with a watermark (WM) indicator according to the availability of a
15 system resource. The ISC, WM and a drop preference (DP) indicator of the data packet are compared to a committed information rate (CIR). The packet is discarded if the DP exceeds the CIR.

Hardware Environment

20 FIG. 1 illustrates a data path of a packet through a hardware environment compatible with the disclosed technology. At step 137, packets 153 enter the L3 card 101 from the access cards 105, 109, trunk card 111, and central process and resource (CPR) 103 cards. At step 139, inbound descriptor information is obtained from the received packet 113 as it is transferred into
25 buffer memory 119 having a free buffer pool 117. The packet type field is compared with PT1 and PT2 type registers, with the PT bits set according. At

step 141, an inbound descriptor 115 is built by combining the CPR 103 and trunk card 111 information. In one embodiment, a channel identifier, control information, and a data packet length is passed directly to the inbound descriptor 115. The CPR 103 descriptor information is forwarded to a FIFO A memory 121 and the trunk card 111 information is forwarded to a FIFO B memory 123. At step 143, the inbound descriptor 125 is read by a route switch processor (CXP) 127 and processed with the help of a command interface 129 and a L3 coprocessor 131. At step 145, an outbound descriptor is written by the CXP 127 to an output queue selector 133. At step 147, the descriptor is directed to the appropriate outbound queue 107. At step 149, the outbound descriptor is used to direct packet re-generation from a DMA controller 135 to a destination 151, and free a buffer in buffer memory 119 back to the free buffer pool 117.

To summarize, the CXP 127 reads an inbound descriptor 125, determines the packet's destination via IP Protocol destination address lookup, builds the packet outbound descriptor for packet transmission, and writes the outbound descriptor to the appropriate channel outbound queue 133 for transmission from a DMA controller 135 to a destination 151.

Buffer Memory

Incoming packets 153 are transferred directly from the access cards 105, 109, trunk card 111, and CPR 103 cards into the L3P buffer memory 119. An example structure of a typical 4MB buffer memory 119, 201 is shown in FIG. 2, where one MB of memory is divided in to 8KB buffers 203, and 3 MB of memory is divided into 2KB buffers 205. A data packet is allocated to a buffer from the 2K or 8K buffer pool as indicated by the length received in the first word of the data packet. Any packet longer than 1536 bytes will be allocated

an 8K byte buffer. If the system is unable to allocate a buffer from the appropriate free buffer pool, the incoming packet is dropped indiscriminately prior to any attempt for proper quality of service (QOS) classification.

Therefore, in order to ensure proper QOS operation, packet discard at this point should be avoided.

The 4MB buffer memory 119, 201, is preferably divided into one or two pools of fixed size buffers, as described above. Two supported sizes are 2KBytes and 8Kbytes, but it will be recognized by one of ordinary skill in the art that other buffer sizes may be used without loss of generality. The division between these two pools is under software control. For example, the entire 4MB can be allocated to all 2KB buffers and vice versa, or different portions can be allocated to each pool.

The pointers to the buffers in the 2KB pool and 8KB pool are maintained in two separate stack memories. Each stack memory can hold up to 8K-1 (or 8191) buffer pointers and typically has a 13-bit stack pointer. The stack pointer always points to the top of the stack. When the Stack Pointer points to the location 0, it indicates that the buffer pool is empty. The content of the stack pointer always represents the number of available buffers in that pool.

To put a buffer pointer on the stack, the stack pointer is advanced by one and then a new buffer pointer is written. To pop a buffer pointer from the stack, the buffer pointer on top of the stack is read and then the stack pointer is subtracted by one.

Watermark Generation

In one embodiment of the disclosed technology, bits 9, 8, and 7 of the 2KB Stack Pointer are used as watermark indicator. These bits provide eight priority buckets as shown in Table 1 below.

5

Watermark	Number of Buffers
0 0 0	0 to 127
0 0 1	128 to 255
0 1 0	256 to 383
0 1 1	384 to 511
1 0 0	512 to 639
1 0 1	640 to 767
1 1 0	768 to 895
1 1 1	> 895

Table 1. Buffer Pool Watermark.

To ensure accurate and timely buffer pool level indication, the watermark bits at the time of the CXP 127 inbound descriptor 125 read are fed into the inbound descriptor 125. The CXP 127 uses these bits to drop packets by freeing up buffers during periods of congestion.

In FIG. 3, an embodiment of the first four words of an inbound descriptor is shown. Word 301 contains a buffer index 303, a sequence number 305, TCP flags 307, a watermark indicator 309, a packet type 311, a bypass indicator 313, a board ID/channel field 315, an error field 317, a control field 319, and a packet length field 321. Words 323, 325, and 327 contain the 20 byte IP header 329. Word 327 also contains a TCP/UDP source port field 331, and a TCP/UDP destination port field 333.

The format and meaning of the watermark indicator 309 is given in Table 1 above. The watermark indicator 309 is used in conjunction with an

internal service class (ISC) and a drop preference bit (DP) to determine when to drop a packet, as shown in Table 2 below.

ISC	DP	Watermark (WM)
7	0	7, 6, 5, 4, 3, 2, 1
7	1	7, 6, 5, 4, 3
6	0	7, 6, 5, 4, 3, 2
6	1	7, 6, 5, 4
5	0	7, 6, 5, 4, 3
5	1	7, 6, 5
4	0	7, 6, 5, 4
4	1	7, 6
3	0	7, 6, 5
3	1	7
2	0	7, 6
2	1	7
1	0	7, 6
1	1	7
0	0	7
0	1	7

Table 2. Congestion Clip Table.

The disclosed technology tags an inbound descriptor 125 with the three bit watermark indicator 309 as they are read by the CXP 127. The three watermark bits of the watermark indicator 309 indicate the current congestion level of the buffer memory 119 and free buffer pool 117. As various thresholds are crossed, the drain rate for the input FIFO A memory 121 and FIFO B memory 123 is increased by discarding lower priority data packets.

Quality of Service (QOS) Classification/Policing

In one embodiment of the disclosed technology, classification and policing are accomplished either implicitly from a type of service (TOS) Octet in the IP header 329, or explicitly as determined by filtering in specific packet

attributes configured by the user. The result of the classification is a 3 bit (8 level) internal service class (ISC) and a drop preference bit (DP). This 4 bit value is used exclusively within the disclosed technology to make QOS related determinations.

5 The policing algorithm shown in FIG. 4 uses user programmable parameters to determine if the traffic rate being received has exceeded the rate of the traffic contract. In certain situations it is desirable to re-classify packets as determined from their traffic utilizations. In these cases the ISC/DP values may be modified during policing operations. At step 401, a lookup key is built
10 using the ISC, DP and watermark values. At 403, a lookup is performed. If at step 405, a threshold is exceeded, at step 407 the number of octet/packet clip drops is counted, and the packet is dropped at step 409. Otherwise, at step 411 the packet is forwarded to the packet destination.

15 QOS Congestion Management

In order to differentiate between contracted service levels there should be traffic management mechanisms in place to ensure that higher precedence traffic has a better chance of making it through the concentrator than best effort traffic, especially in the case of system congestion. This system
20 congestion is indicated by low resources in the free buffer pool 117. It is desirable to ensure 'premium' class traffic over 'best effort'.

In order to make an intelligent choice on which packets to discard and which to keep, at least three parameters are evaluated: 1) the severity of the resource shortage, indicated by watermark bits 309 in the inbound descriptor
25 125, 2) the ISC of the packet, and 3) the DP indication, a traffic rate exceeding its committed information rate (CIR) may be subject for quick discard. The CXP 127 performs this evaluation by concatenating the ISC, DP, and WM

values into a key for a lookup into a configurable congestion clip table, an example of which is shown above in Table 2. The lookup will either 'hit', indicating the packet continues to be forwarded, or 'miss' in which case the packet is discarded or 'clipped'.

5 FIG. 5 shows an example flow diagram of a method for discarding data packets. At step 501, a data packet is classified according to a type of service (TOS) indicator. At step 503, the data packet is modified with an internal service class (ISC) indicator according to the TOS indicator. At step 505, the ISC is compared to a committed information rate (CIR). If at step 507 the ISC
10 exceeds the CIR, the data packet is dropped at step 509.

It is desirable that in these congested scenarios that the lesser precedence packets can be dropped at a rate fast enough to allow for higher precedence packets to continue to be forwarded and also to drain the inbound descriptor queues (FIFO A memory 121 and FIFO B memory 123) in order to
15 replenish the free buffer pool 117.

While the invention is described in terms of preferred embodiments in a specific system environment, those of ordinary skill in the art will recognize that the invention can be practiced, with modification, in other and different hardware and software environments within the spirit and scope of the
20 appended claims.

CLAIMS

What is claimed is:

- 1 1. A method for discarding a data packet comprising the steps of:
2 classifying the data packet according to a type of service (TOS) indicator;
3 modifying the data packet with an internal service class (ISC) indicator
4 according to the TOS indicator;
5 comparing the ISC to a committed information rate (CIR); and
6 discarding the packet if the ISC exceeds the CIR.
- 1 2. The method of claim 1 wherein the step of comparing the ISC to the
2 CIR comprises the steps of:
3 finding an entry in a congestion clip table (CCT) using the ISC as a key
4 value; and
5 comparing the entry to the CIR.
- 1 3. The method of claim 1 wherein the step of classifying the data packet
2 according to a TOS indicator comprises the steps of:
3 analyzing a field of the data packet to determine a packet characteristic;
4 and
5 assigning the TOS indicator based upon the packet characteristic.

1 4. A method for discarding a data packet comprising the steps of:
2 classifying the data packet according to a type of service (TOS) indicator;
3 modifying the data packet with an internal service class (ISC) indicator
4 according to the TOS indicator;
5 modifying the data packet with a watermark (WM) indicator according
6 to the availability of a system resource;
7 comparing the ISC, WM and a drop preference (DP) indicator of the
8 data packet to a committed information rate (CIR); and
9 discarding the packet if the DP exceeds the CIR.

1 5. The method of claim 4 wherein the step of comparing the ISC, WM
2 and DP to the CIR comprises the steps of:
3 concatenating the ISC, WM and the DP into a key value;
4 finding an entry in a congestion clip table (CCT) using the ISC as a key
5 value; and
6 comparing the entry to the CIR.

1 6. The method of claim 5 wherein the step of classifying the data packet
2 according to a TOS indicator comprises the steps of:
3 analyzing a field of the data packet to determine a packet characteristic;
4 and
5 assigning the TOS indicator based upon the packet characteristic.

1 7. An apparatus for discarding a data packet comprising:
2 a classifier to classify the data packet according to a type of service (TOS)
3 indicator;
4 a modifier logically coupled to the classifier to modify the data packet
5 with an internal service class (ISC) indicator according to the TOS indicator;
6 a comparator logically coupled to the modifier to compare the ISC to a
7 committed information rate (CIR); and
8 a discarder logically coupled to the comparator to discard the packet if
9 the ISC exceeds the CIR.

1 8. The apparatus of claim 7 wherein the comparator comprises:
2 a congestion clip table (CCT) having an entry indexed by the ISC; and
3 a comparator logically coupled to the modifier to compare the entry to a
4 committed information rate (CIR).

1 9. The apparatus of claim 7 wherein the classifier comprises:
2 an analyzer to analyze a field of the data packet to determine a packet
3 characteristic; and
4 an assigner logically coupled to the analyzer to assign the TOS indicator
5 based upon the packet characteristic.

1 **10.** An apparatus for discarding a data packet comprising:
2 a classifier to classify the data packet according to a type of service (TOS)
3 indicator;
4 a first modifier logically coupled to the classifier to modify the data
5 packet with an internal service class (ISC) indicator according to the TOS
6 indicator;
7 a second modifier logically coupled to the classifier to modify the data
8 packet with a watermark (WM) indicator according to the availability of a
9 system resource;
10 a comparitor logically coupled to the modifier to compare the ISC, WM
11 and a drop preference (DP) indicator of the data packet to a committed
12 information rate (CIR); and
13 a discarder logically coupled to the comparitor to discard the packet if
14 the DP exceeds the CIR.

1 **11.** The apparatus of claim 7 wherein the comparitor comprises:
2 a concatenator to concatenate the ISC, WM and the DP into a key value;
3 a congestion clip table (CCT) having an entry indexed by the key value;
4 and
5 a comparitor logically coupled to the modifier to compare the entry to a
6 committed information rate (CIR).

1 **12.** The apparatus of claim 7 wherein the classifier comprises:
2 an analyzer to analyze a field of the data packet to determine a packet
3 characteristic; and
4 an assigner logically coupled to the analyzer to assign the TOS indicator
5 based upon the packet characteristic.

1 **13.** An article of manufacture for use in a computer system to discard a
2 data packet, the article of manufacture comprising a computer usable
3 medium having computer readable program code means embodied in the
4 medium, the program code means including:

5 computer readable program code means embodied in the computer
6 usable medium for causing a computer to classify the data packet according to
7 a type of service (TOS) indicator;

8 computer readable program code means embodied in the computer
9 usable medium for causing a computer to modify the data packet with an
10 internal service class (ISC) indicator according to the TOS indicator;

11 computer readable program code means embodied in the computer
12 usable medium for causing a computer to compare the ISC to a committed
13 information rate (CIR); and

14 computer readable program code means embodied in the computer
15 usable medium for causing a computer to discard the packet if the ISC exceeds
16 the CIR.

1 **14.** The article of manufacture of claim 13 wherein the computer readable
2 program code means embodied in the computer usable medium for causing a
3 computer to compare the ISC to the CIR comprises:

4 computer readable program code means embodied in the computer
5 usable medium for causing a computer to find an entry in a congestion clip
6 table (CCT) using the ISC as a key value; and

7 computer readable program code means embodied in the computer
8 usable medium for causing a computer to compare the entry to the CIR.

1 **16.** An article of manufacture for use in a computer system to discard a
2 data packet, the article of manufacture comprising a computer usable
3 medium having computer readable program code means embodied in the
4 medium, the program code means including:

5 computer readable program code means embodied in the computer
6 usable medium for causing a computer to classify the data packet according to
7 a type of service (TOS) indicator;

8 computer readable program code means embodied in the computer
9 usable medium for causing a computer to modify the data packet with an
10 internal service class (ISC) indicator according to the TOS indicator;

11 computer readable program code means embodied in the computer
12 usable medium for causing a computer to modify the data packet with a
13 watermark (WM) indicator according to the availability of a system resource;

14 computer readable program code means embodied in the computer
15 usable medium for causing a computer to compare the ISC, WM and a drop
16 preference (DP) indicator of the data packet to a committed information rate
17 (CIR); and

18 computer readable program code means embodied in the computer
19 usable medium for causing a computer to discard the packet if the DP exceeds
20 the CIR.

17. The article of manufacture of claim 16 wherein the computer readable program code means embodied in the computer usable medium for causing a computer to compare the ISC, WM and DP to the CIR comprises:

computer readable program code means embodied in the computer usable medium for causing a computer to concatenate the ISC, WM and the DP into a key value;

computer readable program code means embodied in the computer usable medium for causing a computer to find an entry in a congestion clip table (CCT) using the key value; and

computer readable program code means embodied in the computer usable medium for causing a computer to compare the entry to the CIR.

18. The article of manufacture of claim 16 wherein the computer readable program code means embodied in the computer usable medium for causing a computer to classify the data packet according to a TOS indicator comprises:

computer readable program code means embodied in the computer usable medium for causing a computer to analyze a field of the data packet to determine a packet characteristic; and

computer readable program code means embodied in the computer usable medium for causing a computer to assign the TOS indicator based upon the packet characteristic.

ABSTRACT

A method, apparatus and article of manufacture for discarding a data packet. The data packet is classified according to a type of service (TOS) indicator and modified with an internal service class (ISC) indicator according to the TOS indicator. The data packet is modified with a watermark (WM) indicator according to the availability of a system resource. The ISC, WM and a drop preference (DP) indicator of the data packet are compared to a committed information rate (CIR). The packet is discarded if the DP exceeds the CIR.

82771.P286




Fig. 1

128
Buffers

1536
Buffers

Buffer Memory ²⁰¹

8KB Buffers
1MB

2KB Buffers
3MB

0x3_FFFF

²⁰³

0x3_0000


0x2_FFFF

²⁰⁵

0x0_0000

128 bits

FIG. 2



first doubleword

63:51	Buffer Index		31:16	Board_id/Channel
50:43	Sequence Number		15	Error
42:38	TCP Flags	A,P,R,S,F (Res, U missing)	14:13	Ctl field
37:35	WM	Watermark	12:0	Length
34:33	PT	Packet Type (see below)		
32	B	Bypass		

second doubleword

63:0 IP Header (bytes 0:7)

third doubleword

63:0 IP Header (bytes 15:8)

fourth doubleword

63:32 IP Header (bytes 19:16)

31:16 UDP/TCP Src Port

PT34:331

04:33]	packet type
0	control packet
1	IPv4 packet
2	IPv6 packet

Fig. 3

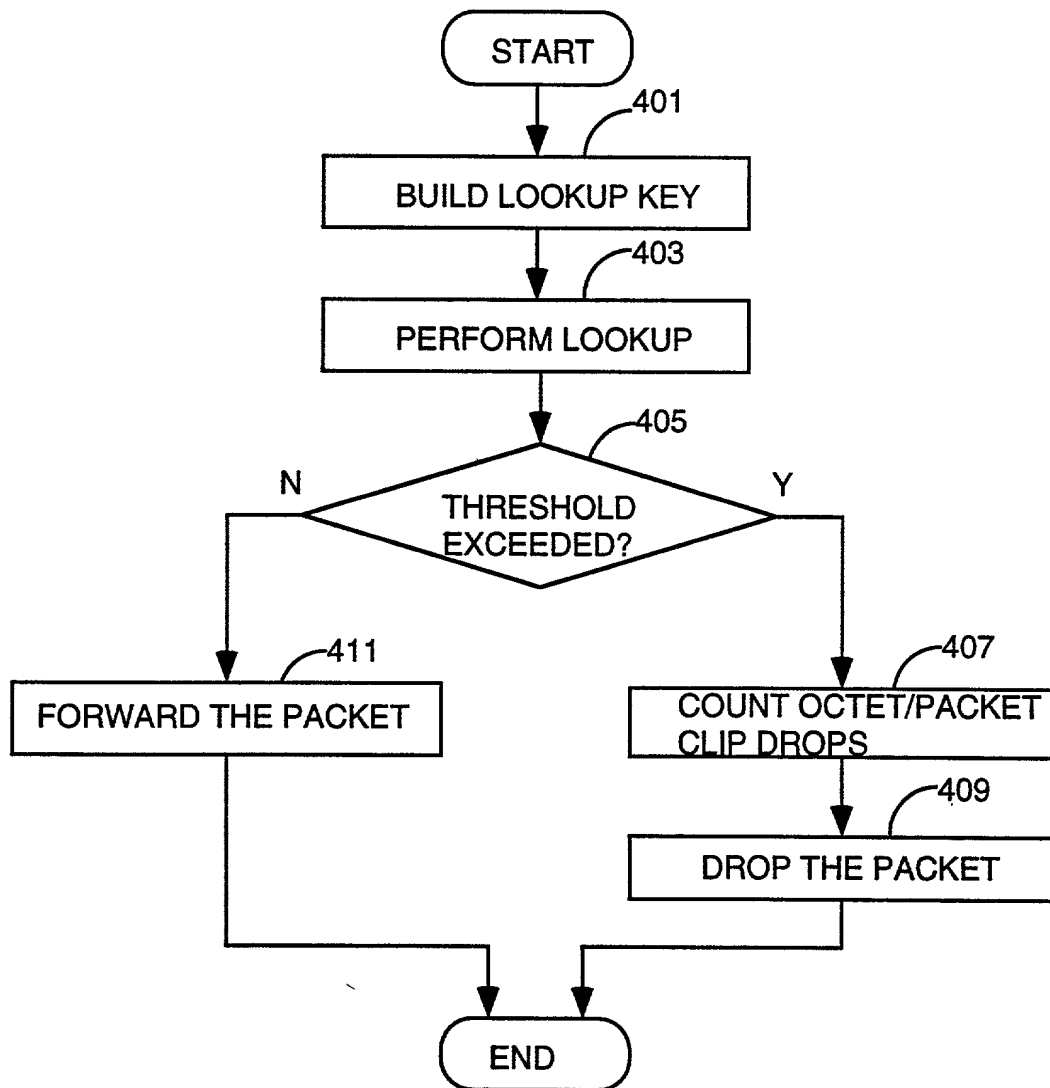


FIG. 4

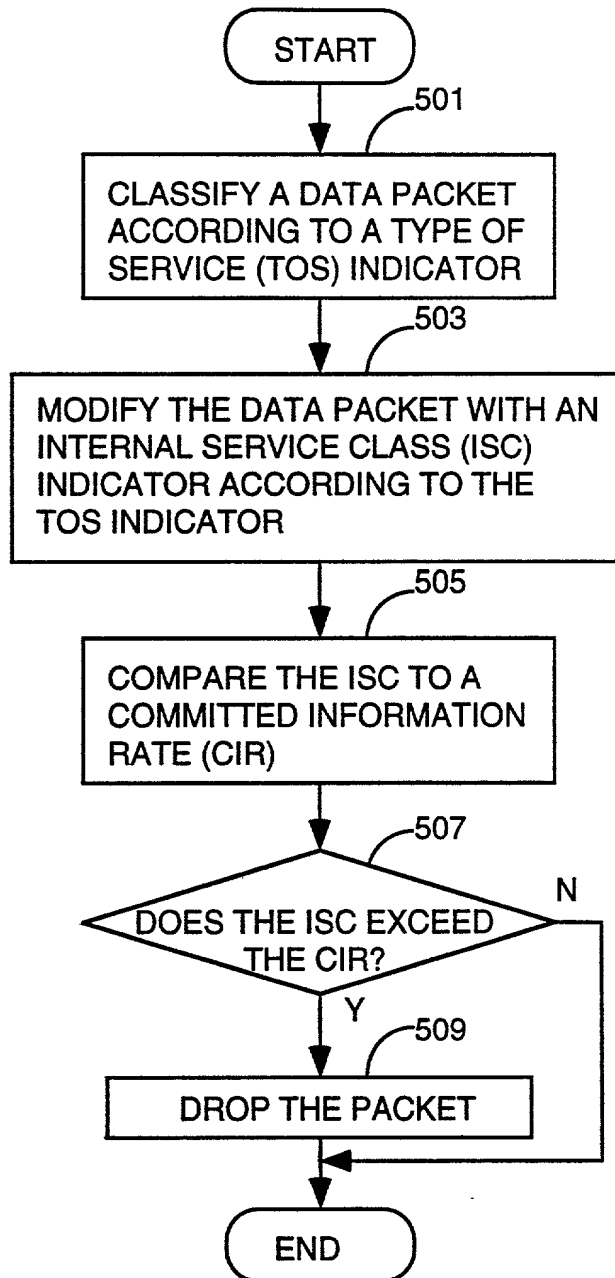


FIG. 5

I hereby claim the benefit under title 35, United States Code, Section 119(e) of any United States provisional application(s) listed below

<u>60/086,078</u>	<u>5/20/98</u>
(Application Number)	Filing Date
<u> </u>	<u> </u>
(Application Number)	Filing Date

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, Section 112, I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

<u> </u>	<u> </u>	<u> </u>
(Application Number)	Filing Date	(Status -- patented, pending, abandoned)
<u> </u>	<u> </u>	<u> </u>
(Application Number)	Filing Date	(Status -- patented, pending, abandoned)

I hereby appoint Farzad E. Amini, Reg. No. P42,261; Aloysius T. C. AuYeung, Reg. No. 35,432; Amy M. Armstrong, Reg. No. 42,265; William Thomas Babbitt, Reg. No. 39,591; Carol F. Barry, Reg. No. 41,600; Jordan Michael Becker, Reg. No. 39,602; Bradley J. Bereznak, Reg. No. 33,474; Michael A. Bernadicou, Reg. No. 35,934; Roger W. Blakely, Jr., Reg. No. 25,831; Gregory D. Caldwell, Reg. No. 39,926; Kent M. Chen, Reg. No. 39,630; Lawrence M. Cho, Reg. No. 39,942; Yong S. Choi, Reg. No. P43,324; Thomas M. Coester, Reg. No. 39,637; Roland B. Cortes, Reg. No. 39,152; Barbara Bokanov Courtney, Reg. No. 42,442; Michael Anthony DeSanctis, Reg. No. 39,957; Daniel M. De Vos, Reg. No. 37,813; Robert Andrew Diehl, Reg. No. 40,992; Tarek N. Fahmi, Reg. No. 41,402; James Y. Go, Reg. No. 40,621; Richard Leon Gregory, Jr., Reg. No. 42,607; Dinu Gruia, Reg. No. P42,996; David R. Halvorson, Reg. No. 33,395; Thomas A. Hassing, Reg. No. 36,159; Phuong-Quan Hoang, Reg. No. 41,839; Willmore F. Holbrow III, Reg. No. P41,845; George W. Hoover II, Reg. No. 32,992; Eric S. Hyman, Reg. No. 30,139; Dag H. Johansen, Reg. No. 36,172; William W. Kidd, Reg. No. 31,772; Michael J. Mallie, Reg. No. 36,591; Andre L. Marais, under 37 C.F.R. § 10.9(b); Paul A. Mendonsa, Reg. No. 42,879; Darren J. Milliken, Reg. No. 42,004; Thinh V. Nguyen, Reg. No. 42,034; Kimberley G. Nobles, Reg. No. 38,255; Michael A. Proksch, Reg. No. 43,021; Babak Redjaian, Reg. No. 42,096; James H. Salter, Reg. No. 35,668; William W. Schaal, Reg. No. 39,018; James C. Scheller, Reg. No. 31,195; Anand Sethuraman, Reg. No. P43,351; Charles E. Shemwell, Reg. No. 40,171; Maria McCormack Sobrino, Reg. No. 31,639; Stanley W. Sokoloff, Reg. No. 25,128; Allan T. Sponseller, Reg. No. 38,318; Judith A. Szepesi, Reg. No. 39,393; Vincent P. Tassinari, Reg. No. 42,179; Edwin H. Taylor, Reg. No. 25,129; George G. C. Tseng, Reg. No. 41,355; Lester J. Vincent, Reg. No. 31,460; John Patrick Ward, Reg. No. 40,216; Stephen Warhola, Reg. No. 43,237; Charles T. J. Weigell, Reg. No. 43,398; Ben J. Yorks, Reg. No. 33,609; and Norman Zafman, Reg. No. 26,250; my attorneys, and James A. Henry, Reg. No. 41,064; Daniel E. Ovanezian, Reg. No. 41,236; Glenn E. Von Tersch, Reg. No. 41,364; and Chad R. Walsh, Reg. No. 43,235; my patent agents, of BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP, with offices located at 12400 Wilshire Boulevard, 7th Floor, Los Angeles, California 90025, telephone (310) 207-3800, and James R. Thein, Reg. No. 31,710, my patent attorney; with full power of substitution and revocation, to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith.

Send correspondence to Thomas A. Hassing, BLAKELY, SOKOLOFF, TAYLOR &
(Name of Attorney or Agent)
ZAFMAN LLP, 12400 Wilshire Boulevard 7th Floor, Los Angeles, California 90025 and direct
telephone calls to Thomas A. Hassing, (408) 720-8598.
(Name of Attorney or Agent)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of Sole/First Inventor Heather Achilles

Inventor's Signature _____ Date _____

Residence _____ (City, State) _____ Citizenship _____ (Country)

Post Office Address _____

Full Name of Second/Joint Inventor Terry Pearson

Inventor's Signature _____ Date _____

Residence _____ (City, State) _____ Citizenship _____ (Country)

Post Office Address _____

Full Name of Third/Joint Inventor Paul Gallo

Inventor's Signature _____ Date _____

Residence _____ (City, State) _____ Citizenship _____ (Country)

Post Office Address _____

Full Name of Fourth/Joint Inventor Larry Manor

Inventor's Signature _____ Date _____

Residence _____ (City, State) _____ Citizenship _____ (Country)

Post Office Address _____

Title 37, Code of Federal Regulations, Section 1.56
Duty to Disclose Information Material to Patentability

(a) A patent by its very nature is affected with a public interest. The public interest is best served, and the most effective patent examination occurs when, at the time an application is being examined, the Office is aware of and evaluates the teachings of all information material to patentability. Each individual associated with the filing and prosecution of a patent application has a duty of candor and good faith in dealing with the Office, which includes a duty to disclose to the Office all information known to that individual to be material to patentability as defined in this section. The duty to disclose information exists with respect to each pending claim until the claim is cancelled or withdrawn from consideration, or the application becomes abandoned. Information material to the patentability of a claim that is cancelled or withdrawn from consideration need not be submitted if the information is not material to the patentability of any claim remaining under consideration in the application. There is no duty to submit information which is not material to the patentability of any existing claim. The duty to disclose all information known to be material to patentability is deemed to be satisfied if all information known to be material to patentability of any claim issued in a patent was cited by the Office or submitted to the Office in the manner prescribed by §§1.97(b)-(d) and 1.98. However, no patent will be granted on an application in connection with which fraud on the Office was practiced or attempted or the duty of disclosure was violated through bad faith or intentional misconduct. The Office encourages applicants to carefully examine:

(1) Prior art cited in search reports of a foreign patent office in a counterpart application, and

(2) The closest information over which individuals associated with the filing or prosecution of a patent application believe any pending claim patentably defines, to make sure that any material information contained therein is disclosed to the Office.

(b) Under this section, information is material to patentability when it is not cumulative to information already of record or being made of record in the application, and

(1) It establishes, by itself or in combination with other information, a prima facie case of unpatentability of a claim; or

(2) It refutes, or is inconsistent with, a position the applicant takes in:

(i) Opposing an argument of unpatentability relied on by the Office, or

(ii) Asserting an argument of patentability.

A prima facie case of unpatentability is established when the information compels a conclusion that a claim is unpatentable under the preponderance of evidence, burden-of-proof standard, giving each term in the claim its broadest reasonable construction consistent with the specification, and before any consideration is given to evidence which may be submitted in an attempt to establish a contrary conclusion of patentability.

(c) Individuals associated with the filing or prosecution of a patent application within the meaning of this section are:

(1) Each inventor named in the application;

(2) Each attorney or agent who prepares or prosecutes the application; and

(3) Every other person who is substantively involved in the preparation or prosecution of the application and who is associated with the inventor, with the assignee or with anyone to whom there is an obligation to assign the application.

(d) Individuals other than the attorney, agent or inventor may comply with this section by disclosing information to the attorney, agent, or inventor.